Volume 385, issue 1 1 November 2007 ISSN 0378-4371

PHYSICA A
STATISTICAL MECHANICS AND ITS APPLICATIONS

Recognized by the European Physical Society

Available online at
ScienceDirect
www.sciencedirect.com

http://www.elsevier.com/locate/physa

Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Physica A 385 (2007) 307–318
PHYSICA A
www.elsevier.com/locate/physa

Noise sensitivity of portfolio selection in constant conditional correlation GARCH models

I. Varga-Haszonits[a,c,*], I. Kondor[a,b]

[a]Department of Physics of Complex Systems, Eötvös University, Pázmány Péter sétány 1/A, H-1117 Budapest, Hungary
[b]Collegium Budapest – Institute for Advanced Study, Szentháromság u. 2, H-1014 Budapest, Hungary
[c]Analytics Department of Fixed Income Division, Morgan Stanley, Budapest, Hungary

Received 24 January 2007; received in revised form 20 May 2007
Available online 26 June 2007

Abstract

This paper investigates the efficiency of minimum variance portfolio optimization for stock price movements following the Constant Conditional Correlation GARCH process proposed by Bollerslev. Simulations show that the quality of portfolio selection can be improved substantially by computing optimal portfolio weights from conditional covariances instead of unconditional ones. Measurement noise can be further reduced by applying some filtering method on the conditional correlation matrix (such as Random Matrix Theory based filtering). As an empirical support for the simulation results, the analysis is also carried out for a time series of S&P500 stock prices.

PACS: 89.65.Gh; 0.545 Tp; 05.40.−a

Keywords: Noisy covariance matrices; Multivariate GARCH models; Constant conditional correlation; Portfolio optimization

1. Introduction

Since the seminal work of Markowitz [1], the problem of finding minimal variance portfolios has been a key issue in financial theory. Mean-variance optimization provides a mathematically clear-cut prescription for obtaining optimal portfolio weights. However, this procedure involves the estimation of the covariance matrix of asset price changes, which may run into severe technical problems even under the assumption that asset returns are "well-behaved" (that is, when they are iid random variables with a multivariate normal distribution).

Denoting the logarithmic change in the price of asset i at time t by x_{it}, and assuming that these changes are independent and identically distributed over time, the sample covariance matrix

$$\widehat{C}_{ij} = \frac{1}{T}\sum_{i=1}^{T}(x_{it} - \overline{x}_i)(x_{jt} - \overline{x}_i) \tag{1}$$

*Corresponding author. Department of Physics of Complex Systems, Etvs University, Pzmny Peter stny 1/A, H-1117 Budapest, Hungary.

E-mail addresses: Istvan.Varga-Haszonits@morganstanley.com (I. Varga-Haszonits), kondor@colbud.hu (I. Kondor).

0378-4371/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.physa.2007.06.017

is a consistent estimator of the process generating covariances C_{ij}, where T is the number of observations (one observation includes N asset price changes), and $\overline{x}_i$ is the mean of x_{it}. However, for short time series, i.e., for values of T/N greater than but close to unity, the fluctuations of the elements $\widehat{C}_{ij}$ over the samples become very large, moreover, for $T/N < 1$ the matrix $\widehat{\mathbf{C}}$ becomes singular which makes portfolio optimization impossible, since it involves matrix inversion.

The effect of measurement noise conspicuously shows up in the eigenvalue spectrum of the sample correlation matrix: Refs. [2,3] found that for realistic values of the ratio T/N a large part of the spectrum can be fitted by the spectral density of a completely random matrix [4], which would seem to imply that most of the information is lost in the noise. Fortunately, the estimation of the covariance matrix and the optimal portfolio weights can be substantially improved by a host of filtering and cleaning procedures, including factor models, Bayesian shrinkage and random matrix theory (RMT) based methods, Refs. [3,5–10]. Filtering methods make the optimization feasible even for $T < N$.

Pafka and Kondor [11] introduced a simulation based method and a measure (denoted by q_0) to quantify the noise sensitivity of the portfolio optimization problem. They found that unless the number of observations exceeds the number of variables significantly (i.e., T is at least 3–5 times N) the sample average of the risk of the estimated portfolio will be at least 20–25% higher than the risk of the true optimum, moreover, in the $T \to N$ limit the portfolio optimization algorithm exhibits critical behavior [12], and the excess risk due to measurement noise grows to infinity. The same methodology can also be used to measure how effectively noise filtering can improve optimization.

However, the scope of the studies mentioned above is usually limited to iid returns, with multivariate normal distribution. Thus, for better practical applicability, these findings have to be generalized by relaxing one or other of the assumptions, and investigating the effects of measurement noise under circumstances that are closer to market realities.

It is well known that daily stock returns are not iid: while their autocorrelation is usually not significant, that of their squares is positive, i.e., price volatility exhibits some degree of persistence. Univariate GARCH models [13,14] are designed to capture this property in the case of individual price movements; however, for the description of the joint stochastic return movements, a more complex model is needed. For this purpose, Bollerslev et al. proposed a general multivariate GARCH model [15], featuring conditional variances for each asset and covariances for each pair of assets. A serious drawback of this framework is that it contains so many parameters that model calibration becomes impossible even for a relatively small number of variables.

In order to overcome this problem, restrictions have to be imposed on the general multivariate GARCH specification, thereby reducing the number of parameters to be estimated. There exist several different versions in the literature (e.g. BEKK model [16], factor-GARCH [17], orthogonal GARCH [18–20], constant conditional correlation (CCC) GARCH [21], time varying and dynamical conditional correlation GARCH [22–24], for a comprehensive survey see Ref. [25]). The one with the simplest specification is perhaps CCC-GARCH, therefore, incorporating it into the portfolio optimization framework represents a reasonable first step towards the generalization of the studies and results on covariance matrix estimation and filtering.

The primary aim of this study is to generalize results on the noise sensitivity of portfolio optimization to a class of non-stationary processes. Throughout this paper, the simulation-based approach of Ref. [11] is followed to analyze the performance of covariance/correlation estimation and portfolio optimization based on noisy data. CCC-GARCH models with different constant correlation matrices are used to simulate the time evolution of stock returns. Optimal portfolio weights are calculated from the generated time series, using the unconditional sample covariance matrix and the conditional covariance matrix computed from CCC-GARCH estimation, respectively. The latter, correctly specified method is expected to outperform the former, which is, indeed, borne out by our study.

The rest of the paper is organized as follows. Section 2 explains how one can characterize the effect of noise on the quality of portfolio selection. Section 3 summarizes how CCC multivariate GARCH models are specified, and how the effect of noise can be measured for CCC-GARCH processes. In Section 4 simulations are used to investigate how the quality of portfolio selection depends on the length of the time series, the process generating model and the way covariances are estimated. In Section 5 empirical data are used to underpin the findings of Section 4.

2. Noise measurement

In order to fix notation, let us recall Markowitź [1] mean-variance optimization scheme. In its simplest setting, it corresponds to the following constrained minimization problem:

$$\min_{\mathbf{w}} \mathbf{w}^{\mathrm{T}}\mathbf{C}\mathbf{w}, \tag{2}$$

$$\mathbf{1}^{\mathrm{T}}\mathbf{w} = 1, \tag{3}$$

$$\mathbf{w}^{\mathrm{T}}E(\mathbf{y}) = \mu, \tag{4}$$

where $\mathbf{w}$ is the vector of portfolio weights, $\mathbf{C}$ is the (conditional or unconditional) covariance matrix of asset returns, the asset return vector is denoted by $\mathbf{y}$, the expected return of the optimal portfolio is μ, and $\mathbf{1}$ is a vector with all components equal to one. Eq. (3) represents the budget constraint, while (4) determines the level of expected return one wishes to realize. The use of the variance as the objective function of the optimization problem implicitly assumes that asset returns are subject to a (conditionally or unconditionally) normal, or some similarly narrow, distribution. Note that, in order to simplify matters and bring out the essential features, we have not imposed any constraint on short selling in the above. In a similar spirit, we also wish to drop the condition (4) on expected return, and restrict our interest to finding the minimum risk portfolio under the budget constraint (3).[1] Then the optimal portfolio weights can be expressed analytically as

$$\mathbf{w} = \frac{\mathbf{C}^{-1}\mathbf{1}}{\mathbf{1}^{\mathrm{T}}\mathbf{C}^{-1}\mathbf{1}}. \tag{5}$$

In a practical situation the "true", or process generating covariance matrix $\mathbf{C}$ is not known, all one can have is the estimator $\widehat{\mathbf{C}}$. At this point, it is not important how this estimator is obtained, whether it is the raw sample covariance matrix or the result of some noise filtering procedure. Since the estimator $\widehat{\mathbf{C}}$ is different from the true covariance matrix $\mathbf{C}$, the optimal portfolio weights $\widehat{\mathbf{w}}$ computed from the former will be different from the $\mathbf{w}$ computed from the latter. The minimal variance is $\mathbf{w}^{\mathrm{T}}\mathbf{C}\mathbf{w}$, but the best weights we can achieve in a finite sample are $\widehat{\mathbf{w}}$, whereas the true process corresponds to $\mathbf{C}$. Thus the measure of suboptimality will be the variance-like quantity $\widehat{\mathbf{w}}^{\mathrm{T}}\mathbf{C}\widehat{\mathbf{w}}$. The ratio of this to the variance of the true optimum has been introduced in Ref. [11] as a measure of the error we make by relying on the noisy data:

$$q_0 = \sqrt{\frac{\widehat{\mathbf{w}}^{\mathrm{T}}\mathbf{C}\widehat{\mathbf{w}}}{\mathbf{w}^{\mathrm{T}}\mathbf{C}\mathbf{w}}}, \tag{6}$$

and will be used for this purpose also in this paper. The measure q_0 is obviously always greater than unity,[2] and it measures exactly what the investor is concerned about: how much excess risk one has to take due to the fact that the true covariance matrix is not known, only its estimator. If q_0 is about 1.1, then the investor is exposed to 1.1 times the risk at the true optimum, which is usually acceptable. However, if q_0 is around 3–4 (which is not an extreme case at all, as we shall see), then the result of the optimization is useless, because we take a risk 3–4 times higher than the real optimum.

The definition of the quantity q_0 refers to a given sample, so it is a random variable. It was shown [11,26] that for $N, T \to \infty$ such that their ratio $r = N/T < 1$ is fixed, the average of q_0 over the samples behaves asymptotically as

$$E(q_0) \approx \frac{1}{\sqrt{1-r}}. \tag{7}$$

The equality is exact if the returns are iid normal, i.e., when $\mathbf{C}$ is the identity matrix, but it also remains true (up to corrections of $O(1/N)$) within a large domain of models containing some of the most popular

[1] Estimating expected returns is a very delicate matter and the additional noise induced by expected return measurement would mask the effect of noise on the covariances.

[2] The notation is taken over from Ref. [11], where quantities denoted by q_1 and q_2 also appear. In the present context the subscript 0 is superfluous, but we keep it for continuity of notation.

correlation structures considered in financial theory (see later in this paper). The message of this strikingly simple result is that for a fixed portfolio size N and for very long observation times T we can get arbitrarily close to the true optimum, but if the number of observations does not sufficiently exceed the number of assets, the estimated optimum will be very far from the true one, moreover, as T approaches N from above, the error diverges! When the number of observations falls below the portfolio size, $T<N$, the empirical covariance matrix $\widehat{\mathbf{C}}$ becomes singular, which makes portfolio optimization impossible. (It is important to keep in mind that (7) is only true when the sample covariance matrix is used to compute the portfolio weights. The divergence at $T=N$ can be suppressed by filtering the covariance matrix [11].)

The divergence of the estimation error has been recognized by Kondor et al. [12] to be the manifestation of an algorithmic phase transition [27] recently. This phase transition is accompanied by a number of related critical phenomena, such as the even stronger divergence of the standard deviation of q_0 [28]. The phase transition context also explains the universality of the asymptotic results, including that of the critical exponent $-1/2$ of $E(q_0)$, which will be confirmed also for the class of non-stationary models to be considered in this paper. Our primary concern here will not be the study of this phase transition, however, we will rather be concerned with how and to what extent its consequences can be remedied.

3. Econometric methodology

In this paper we investigate the effectiveness of portfolio optimization, with or without the filtering described above, for stock returns following a multivariate GARCH process, namely the CCC-GARCH model of Bollerslev [21]. This choice is motivated by the simplicity and easy specification of the model, leaving open the possibility of investigating more complicated models after the CCC-GARCH case is satisfactorily understood. Before introducing the CCC-GARCH process, we briefly recall its univariate predecessor, the GARCH process [14].

3.1. The GARCH process

Processes conforming the following specification are called GARCH(1,1):

$$y_t = \mu + \varepsilon_t, \tag{8}$$

$$\varepsilon_t = \sqrt{h_t} v_t, \tag{9}$$

$$h_t = \gamma + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}, \tag{10}$$

$$\gamma, \alpha, \beta \geqslant 0,$$

where v_t is an iid standard normal process and γ, α, β and μ are constants and $\alpha+\beta\leqslant 1$. For empirical time series these four parameters can be estimated using the maximum likelihood method. (8) and (9) together mean that y_t is normal with a mean μ and time dependent variance h_t. The time dependence of the variance is determined by (10), which expresses that the variance at time period t is determined by the variance of y and the size of the movement of y over the preceding time period $t-1$. Since h_t can be predicted based on information available at $t-1$, h_t is called the conditional variance of y_t.

The most important property of this process due to (10) is that the conditional variance exhibits persistence leading to the so called "volatility clustering" effect which is a well-known feature of empirical stock price movements.[3]

It is worth noting that y_t is only conditionally normal, based on information available at $t-1$. In contrast, the unconditional distribution of y is fat-tailed (capturing another important feature of stock price movement). The higher the value of $\alpha+\beta$, the higher the (unconditional) variance and fatter the tails. When $\alpha+\beta=1$ the unconditional variance becomes infinite.

It should also be mentioned that GARCH(1,1) is a special case of the GARCH(p,q) processes where the conditional variance equation (10) includes more than one time-lagged values of ε_t and h_t (up to ε_{t-q} and h_{t-p}).

[3]"Volatility clustering" refers to the phenomenon that the random process governing stock returns is not stationary: periods of higher and lower variance follow one another forming visible clusters when the stock returns are plotted against time.

3.2. The CCC-GARCH(1,1) process

CCC-GARCH(1,1) [21] is a straightforward generalization of GARCH(1,1):

$$y_{i,t} = \mu_i + \varepsilon_{i,t}, \tag{11}$$

$$h_{ii,t} = \gamma_i + \alpha_i \varepsilon_{i,t-1}^2 + \beta_i h_{ii,t-1}, \tag{12}$$

$$h_{ij,t} = \rho_{ij}\sqrt{h_{ii,t}h_{jj,t}},$$

$$\gamma_i, \alpha_i, \beta_i \geqslant 0, \quad i = 1, 2, \ldots, N, \tag{13}$$

where the parameters to be estimated are μ_i (expected returns), γ_i, α_i, β_i and the ρ_{ij} (CCC). Eq. (11) expresses that the expected return on stock i is μ_i, while (12) models the time evolution of the variance of price changes for stock i. The time varying conditional covariance matrix is calculated from the constant correlations and the time varying variances according to (13).

This model can be calibrated in two steps: first a univariate GARCH(1,1) model is fitted to the individual stock price movements, then the conditional correlation matrix is calculated from the GARCH residuals ($\widehat{\rho}_{ij}$ itself is, actually, not a correlation matrix, but a consistent estimator of ρ_{ij}):

$$v_i = \frac{y_{it} - \widehat{y}_{it}}{\sqrt{\widehat{h}_{ii,t}}}, \tag{14}$$

$$\widehat{\rho}_{ij} = \frac{1}{T}\sum_{i=1}^{T} v_{it}v_{jt}, \tag{15}$$

$$\widehat{h}_{ij,t} = \widehat{\rho}_{ij}\sqrt{\widehat{h}_{ii,t}\widehat{h}_{jj,t}}. \tag{16}$$

In the present analysis we simulate time series using the model specified above, and investigate how q_0 depends on N, T and ρ_{ij}. However, the precise definition of q_0 for GARCH models needs some further discussion. Both the true conditional covariance matrix and the estimated one are time dependent now. For a given time t the true conditional covariance matrix is $\mathbf{H}_t$, but within a finite sample, we have only access to the estimate $\widehat{\mathbf{H}}_t$. If the optimal portfolio weights corresponding to these two matrices are $\mathbf{w}$ and $\widehat{\mathbf{w}}$, respectively, we can define the following quantity:

$$q_0^c = \sqrt{\frac{\widehat{\mathbf{w}}^{\mathrm{T}}\mathbf{H}_t\widehat{\mathbf{w}}}{\mathbf{w}^{\mathrm{T}}\mathbf{H}_t\mathbf{w}}}, \tag{17}$$

where the superscript c refers to the fact that the optimal portfolio weights are calculated from the *conditional* covariance matrix. It is also interesting to see what happens when the asset yields are treated as if they were simply iid, and the optimal portfolio is computed from the *unconditional* sample covariance matrix:

$$\widehat{G}_{ij} = \frac{1}{T}\sum_{t=1}^{\mathrm{T}}(y_{it} - \overline{y}_i)(y_{jt} - \overline{y}_j), \tag{18}$$

where $\overline{y}_i$ is simply the mean of y_{it} along the time series. Let $\widehat{\mathbf{u}}$ denote the optimal portfolio weights calculated from $\widehat{\mathbf{G}}$, then we can define

$$q_0^u = \sqrt{\frac{\widehat{\mathbf{u}}^{\mathrm{T}}\mathbf{H}_t\widehat{\mathbf{u}}}{\mathbf{w}^{\mathrm{T}}\mathbf{H}_t\mathbf{w}}}. \tag{19}$$

Both measures will be calculated for every simulation, which will allow us to draw some conclusions regarding the extent to which GARCH estimation can improve portfolio selection with respect to estimates based on raw data.

4. Simulation results

We have carried out simulations to measure the noise sensitivity of portfolio selection with stock prices following a CCC-GARCH(1,1) process. We also tested how noise reduction techniques applied to the conditional and unconditional covariance matrices can affect the performance of optimization when combined with CCC-GARCH(1,1) estimation. For this purpose, we chose the RMT-based method introduced in Refs. [9,29], which is fairly simple and has already been discussed in the context of multivariate GARCH models in Ref. [30]. Other methods (e.g. factor models, Bayesian shrinkage) might also be employed, but that would not change the main conclusions of this study.

We have simulated multivariate CCC-GARCH time series with different CCC. Four different correlation models were chosen (similarly to those used in Ref. [11]):

- Model I: The stock price movements are uncorrelated, the correlation matrix is the identity.
- Model II: Asset prices are uniformly correlated, the correlations are the same (ρ_0) for any pair of stocks. (This is basically a one factor model.)
- Model III: Stocks can be divided into a number of sectors. The size of the sectors are the same, the correlation between two stocks within the same sector is ρ_1, the correlation for two stocks belonging to different sectors is ρ_0, where $\rho_0 < \rho_1$. The correlation matrix is block diagonal, with equally sized blocks along the main diagonal.
- Model IV: Sector model with sectors of different sizes, and correlations. The correlation between two stocks belonging to separate sectors is ρ_0, while the correlation within sector i is $\rho_i > \rho_0$. The corresponding correlation matrix is again block diagonal, but the blocks are different in size and correlation.

Note that Models III and IV fit into the correlation modeling framework proposed by Noh [31].

The number of simulated assets was $N = 100$, and two sets of GARCH parameters were tested. In both cases, stock returns were assumed to follow the same univariate GARCH(1,1) process, with conditional mean $\mu_i = 0$ for each asset. In the first case, the GARCH parameters were chosen to be $\gamma_i = 0.03$. $\alpha_i = 0.1$ and $\beta_i = 0.6$ (for $i = 1, \ldots, N$). In the second case the parameters were $\gamma_i = 0.03$. $\alpha_i = 0.15$ and $\beta_i = 0.8$ (for $i = 1, \ldots, N$). For Model II $\rho_0 = 0.3$, for Model III $\rho_0 = 0.2$ and $\rho_1 = 0.6$, while for model IV we chose $\rho_0 = 0.2$, and we set five sectors containing 30, 10, 20, 15 and 25 stocks, respectively, with intra-sector correlations of 0.5, 0.55, 0.6, 0.65 and 0.7.

For each correlation model and both sets of GARCH parameters (i.e., for eight different specifications) time series of length $T = 110$, 150, 200, 300, 400, 600, 800, 1200, 1600, 3200 and 6400 were simulated. These time series were then used to estimate GARCH parameters and CCC, and optimal portfolio weights were calculated according to four different methods:

- Method I: optimization based on the unconditional covariance matrix.
- Method II: optimization based on the unconditional covariance matrix, with RMT filtering.
- Method III: optimization based on the conditional covariance matrix.
- Method IV: optimization based on the conditional covariance matrix, with RMT filtering.

The first two methods treat the data as if they were normally distributed iid time series, and, accordingly, the performance of the optimal portfolio is measured by q_0^u. Methods III and IV use the estimated conditional covariance matrix for portfolio selection, so the performance is measured by q_0^c (using the RMT filtered correlation matrix for Methods II and IV). Every simulation was repeated 20 times, and the mean of the resulting q_0 values was computed.

For the first set of GARCH parameters results for pre-filtering optimization (Methods I and III) can be seen in Fig. 1. In order to see deviations from Eq. (7), q_0^u and q_0^c are plotted against $1 - N/T$, on a log–log scale, along with the straight line corresponding to (7). For Model I, there is no significant difference between the r dependence of q_0^u and q_0^c, while for the other models, the use of conditional variances yields a slight improvement in the value of q_0: in all of the graphs for Models II–IV, q_0^u is systematically higher than q_0^c (see also Fig. 2 on a linear scale).

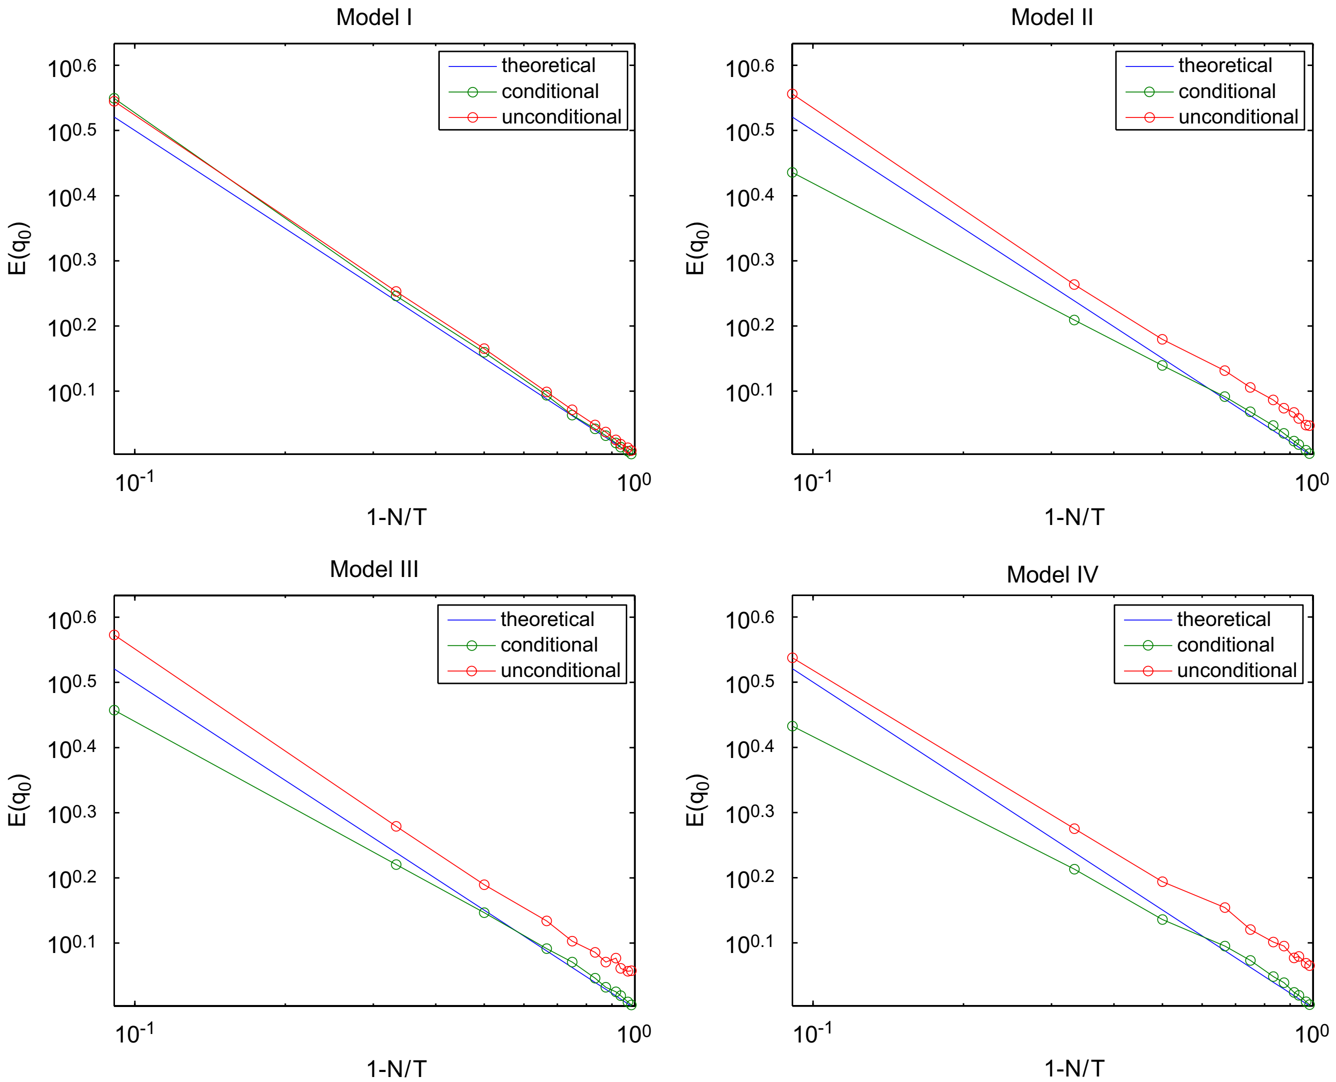

Fig. 1. Dependence of $E(q_0)$ on the length of the time series for different correlation models (log–log plot). Process generating GARCH parameters: $\gamma_i = 0.03$, $\alpha_i = 0.1$, $\beta_i = 0.6$ for every $i = 1, \ldots, N$. The "theoretical" curve corresponds to Eq. (7), the "conditional" curve to q_0^c and the "unconditional" curve to q_0^u.

For the second set of GARCH parameters the sum $\alpha + \beta$ is higher than in the first case, and it is close to one. This implies that the tails of the unconditional distribution are fatter, and conditional variances are more persistent. Therefore, one expects that the optimal portfolio weights estimated from the unconditional covariances will contain a higher level of measurement error. In fact, Fig. 3 shows that for Models II–IV q_0^u is substantially higher than q_0^c, especially for short time series (see also Fig. 4). For longer time series, $E(q_0^u)$ stabilizes around 2, and, at least within the range of our simulations, it does not seem to significantly decrease any further (see Fig. 5). This leads to the conclusion that in the presence of strong GARCH effects (i.e., where $\alpha + \beta$ is close to unity), it is essential to use conditional covariances for portfolio selection, otherwise the volatility of our chosen portfolio can substantially exceed its optimal value.

Another important result is that $E(q_0^c)$ follows the relation (7), at least in the limited range studied, quite closely for each Model, which is a manifestation of universality.[4] Thus, if the stock prices follow a CCC-GARCH process, and it is known that it is a CCC-GARCH process (i.e., our model specification is

[4]In order to have a more stringent test of universality, we should, of course, run the simulations in a much wider range of the critical parameter $1 - N/T$, but this numerically demanding task is beyond the scope of the present work. It would also be interesting to compare the higher moments of q_0 in the various cases, not only its average.

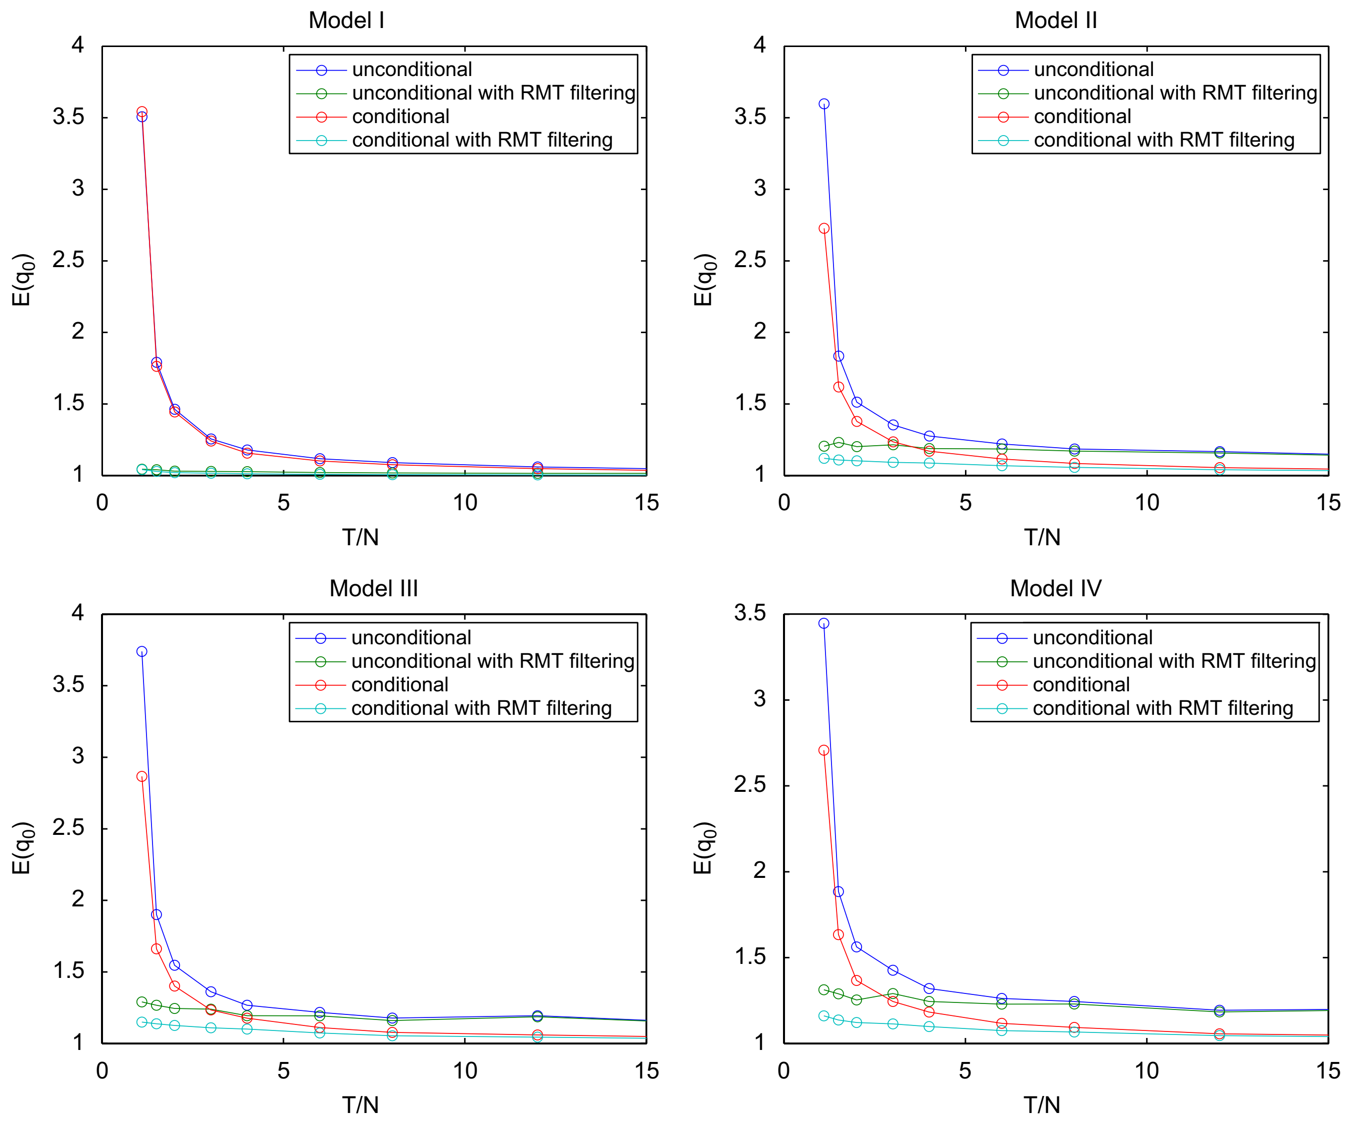

Fig. 2. Effect of RMT filtering for different correlation models. Process generating GARCH parameters: $\gamma_i = 0.03$, $\alpha_i = 0.1$, $\beta_i = 0.6$ for every $i = 1, \ldots, N$. The "conditional" curve corresponds to q_0^c and the "unconditional" curve to q_0^u, while "conditional with RMT filtering" and "unconditional with RMT filtering" refer to Methods II and IV.

correct), then the increase in risk due to measurement noise can be reduced to a level very close to what we would experience in the case of iid normal stock price fluctuations (discussed extensively in Refs. [11,32,33]).

This approximation becomes exact in the "thermodynamic limit" (i.e., when $N, T \to \infty$ with $N/T = \text{const}$). On the one hand, the GARCH estimation is consistent, so the statistical error of estimated GARCH parameters and residuals will converge to zero as the number of observations T goes to infinity, independently of the number of variables N. On the other hand, the estimation error of the correlations depends only on the ratio N/T, hence, in the thermodynamic limit the error of the GARCH parameters will disappear, while that of the conditional correlations will converge to some non-zero value.[5] Consequently, in this limit the only source of noise will be the constant correlation matrix estimated from the GARCH residuals, which are themselves iid normal.

Furthermore, whether we use conditional or unconditional covariances, RMT filtering can reduce noise very efficiently, leading to a better portfolio and lower q_0, independently of the specific parameters of the

[5]More precisely, to some well-defined distribution with positive mean and standard deviation.

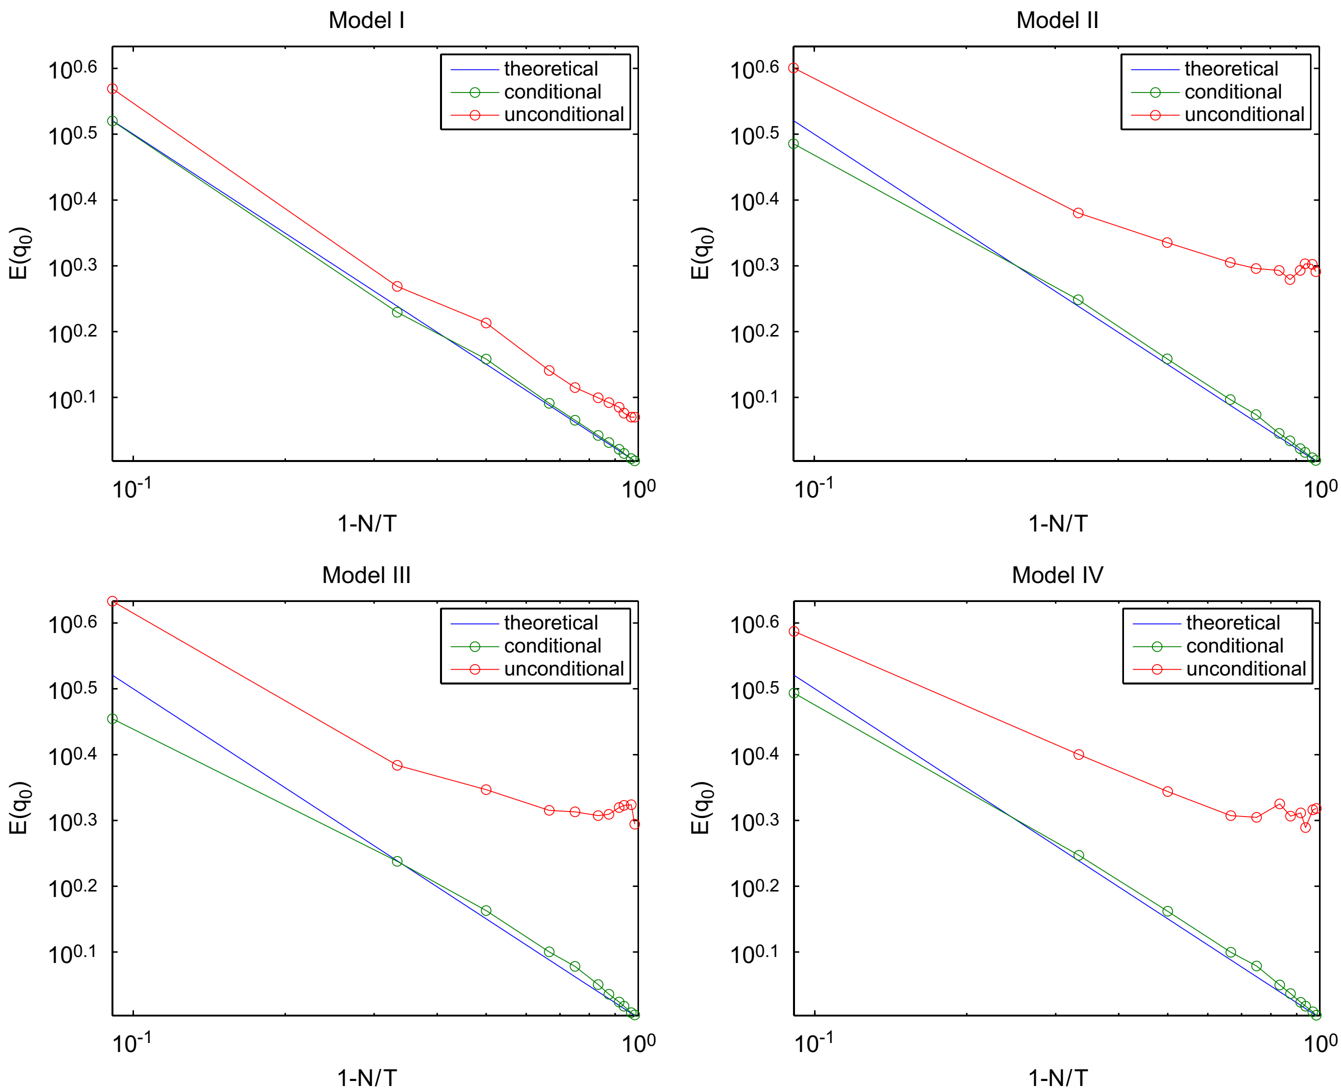

Fig. 3. Dependence of $E(q_0)$ on the length of the time series for different correlation models (log–log plot). Process generating GARCH parameters: $\gamma_i = 0.03$, $\alpha_i = 0.15$, $\beta_i = 0.8$ for every $i = 1, \ldots, N$, the conditional correlation matrix is that of Model I. The "theoretical" curve corresponds to Eq. (7), the "conditional" curve to q_0^c and the "unconditional" curve to q_0^u.

underlying process (see Figs. 2 and 4). The improvement is particularly significant when T/N is of order unity. (For large T/N values the sample covariance matrices are less noisy, and the resulting optimal portfolios are already reasonably good without RMT filtering.) After RMT filtering q_0 does not diverge any more, moreover, it is also possible to optimize portfolios in the $T/N < 1$ range, with acceptably low q_0 values (see Ref. [11]).

It may be surprising—especially for the second set of GARCH parameters—that filtering the unconditional covariance matrix provides a better estimate for shorter time series than for longer ones, while intuitively one would expect that the longer time series result in lower q_0 even for filtered covariance matrices (as it is the case for conditional optimization). The most plausible explanation for this phenomenon is the following. Unconditional optimization is based on a misspecified model (namely that the variances and covariances are constant, and price changes are iid). For small T, the time series may not be long enough for the GARCH effect to manifest itself, thus the error that we make by treating it as iid is relatively small. This error, unlike that due to measurement noise, increases with T. Thus, although RMT filtering removes a great part of the random noise due to the finiteness of the sample, it cannot deal with errors resulting from the misspecification of the estimated model. The local minima in Fig. 5 can also be explained

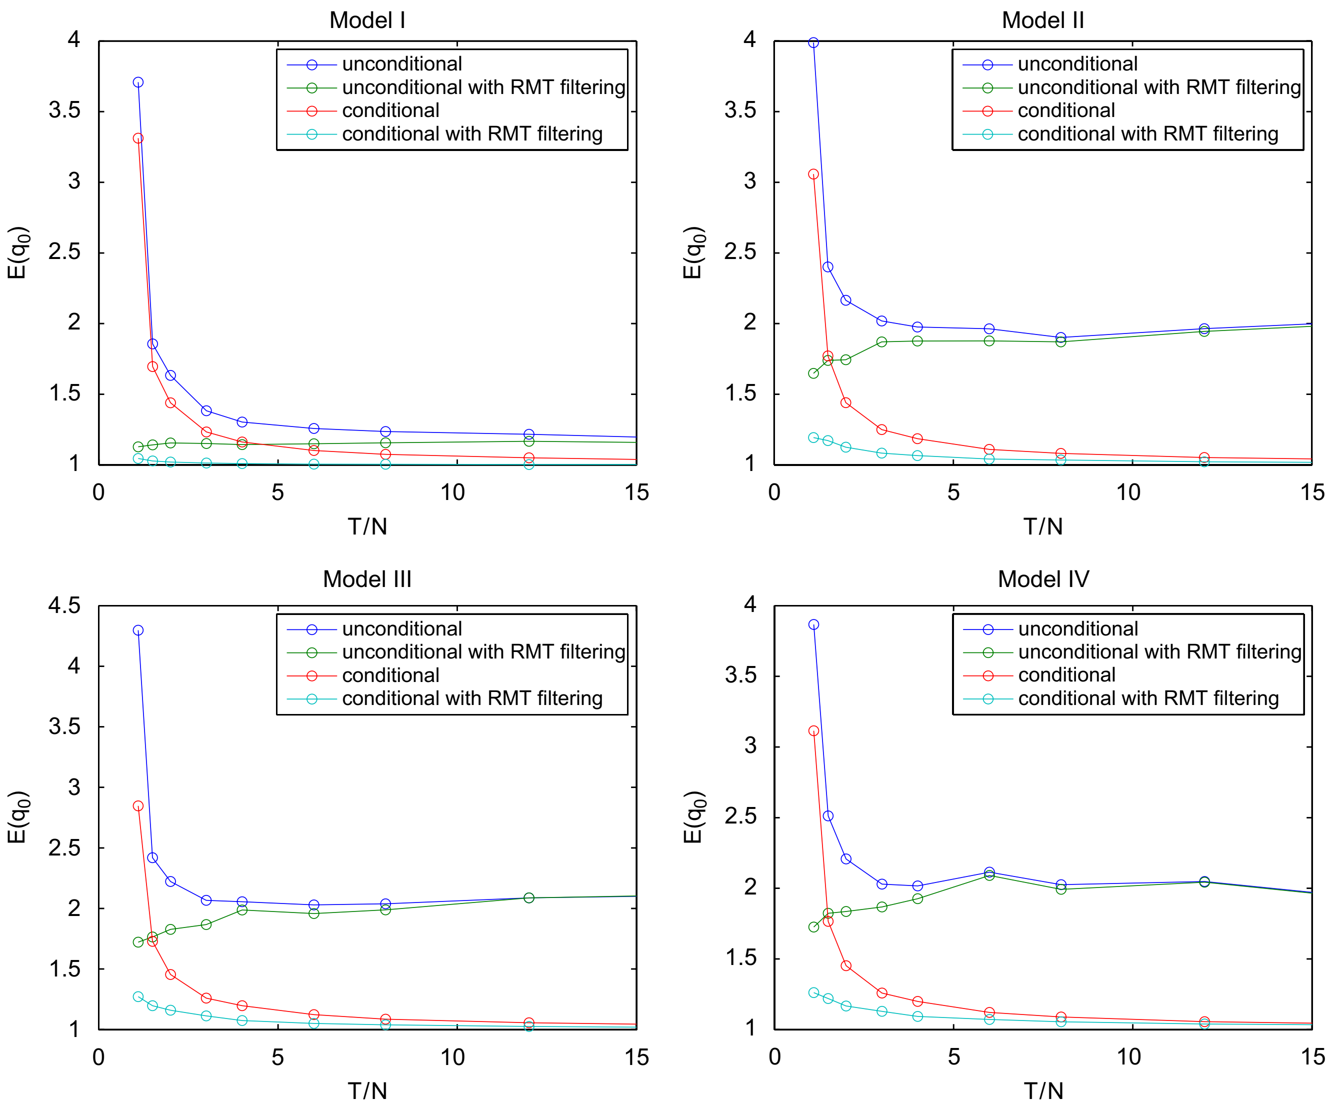

Fig. 4. Effect of RMT filtering for different correlation models. Process generating GARCH parameters: $\gamma_i = 0.03$, $\alpha_i = 0.15$, $\beta_i = 0.8$ for every $i = 1, \ldots, N$, the conditional correlation matrix is that of Model I. The "conditional" curve corresponds to q_0^c and the "unconditional" curve to q_0^u, while "conditional with RMT filtering" and "unconditional with RMT filtering" refer to Methods II and IV.

by this effect: they occur where the total impact of these two types of error on the selected portfolio is minimal.

5. Empirical results

In order to test the above methods on real market data, we also performed an empirical analysis (similar to the study presented in Ref. [30], but with a different set of data) to test the effectiveness of portfolio optimization based on CCC-GARCH estimation. We used $N = 406$ S&P500 stock prices from the period 1991 to 1996, computed daily logarithmic changes (the number of observations is $T = 1308$), and used Methods I–IV to estimate the minimal variance portfolios.

For Methods I and II the first 812 observations (so that $r = N/T = 0.5$) were used to calculate the covariance matrix and the optimal portfolio weights, as well as the volatility that would have been realized, if this portfolio had been held for the remaining 496 business days.

For Methods III and IV a moving window approach was used. For a given business day, the conditional covariance matrix was calculated by estimating a multivariate CCC-GARCH model based on the preceding 812 business days. Thus, the optimal portfolio corresponding to the conditional covariance matrix was

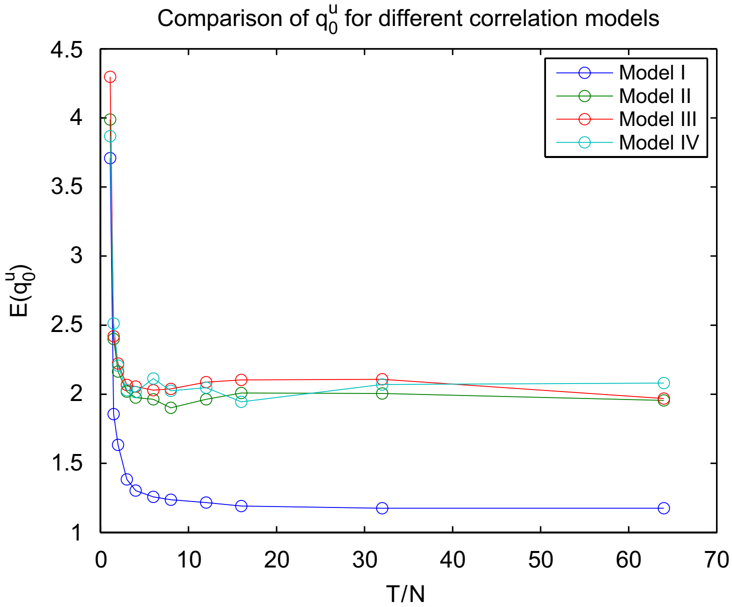

Fig. 5. Comparison of the T-dependence of q_0^u for different correlation models. Process generating GARCH parameters: $\gamma_i = 0.03$, $\alpha_i = 0.15$, $\beta_i = 0.8$ for every $i = 1, \ldots, N$. For Models II, III (around $T = 8$) and IV (around $T = 16$) shallow local minima are observable.

readjusted on a daily basis, and the volatility realized over the last 496 business days was computed again. The results are summarized in the following table:

Method	Volatility (%)
Uniform	6.13×10^{-3}
I	5.31×10^{-3}
II	4.32×10^{-3}
III	4.77×10^{-3}
IV	3.88×10^{-3}

For comparison, the daily volatility of the uniformly weighted portfolio is also shown. Uniform weighting yields the highest volatility, while the other four are ranked as IV, II, III, I, with Method IV performing best. This empirical example also supports our conclusion, that the quality of portfolio selection can be substantially improved by multivariate GARCH modeling.

6. Conclusion

We have investigated the efficiency of minimal variance portfolio estimation, when the underlying process is a Constant Conditional Correlation GARCH. CCC-GARCH(1,1) processes with different parameters and conditional correlation matrices were simulated, and the optimal portfolios were calculated both for the correct CCC-GARCH model specification and for the incorrect, unconditional covariance matrix specification. Both specifications were used with and without random matrix filtering. Unconditional portfolio optimization was found to perform substantially worse than conditional optimization, especially when the sum of GARCH parameters was close to 1. Random matrix filtering was found to significantly improve the quality of portfolio selection, but this improvement can be offset by errors due to misspecification. The methods in consideration were also tested on empirical data, which confirmed the simulation-based results. The main conclusion of this study is that GARCH effects present in asset price movements should be identified, and portfolio selection should be performed on the basis of the conditional covariance matrix. For GARCH processes the efficiency of portfolio selection will be the same as in the case of iid processes, provided that we use a correctly specified GARCH model to describe the data. In contrast, neglecting conditional heteroscedasticity may considerably increase the risk of the estimated optimal portfolio.

Acknowledgment

This work has been supported by the "Cooperative Center for Communication Networks Data Analysis", a NAP project sponsored by the National Office of Research and Technology under Grant no. KCKHA005.

References

[1] H. Markowitz, Portfolio Selection: Efficient Diversification of Investments, Wiley, New York, 1959.
[2] L. Laloux, P. Cizeau, J.-P. Bouchaud, M. Potters, Noise dressing of financial correlation matrices, Phys. Rev. Lett. 83 (1999) 1467.
[3] V. Plerou, P. Gopikrishnan, B. Rosenow, L.A.N. Amaral, H.E. Stanley, Universal and non-universal properties of cross-correlations in financial time series, Phys. Rev. Lett. 83 (1999) 1471.
[4] V.A. Marchenko, L.A. Pastur, Math. USSR-*Sb* (1967) 457.
[5] O. Ledoit, M. Wolf, A well-conditioned estimator for large-dimensional covariance matrices, J. Multivar. Anal. 88 (2) (2004) 365–411.
[6] O. Ledoit, M. Wolf, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, J. Empirical Finance 10 (5) (2003) 603–621.
[7] O. Ledoit, M. Wolf, Honey, I shrunk the sample covariance matrix, J. Portfolio Manage. 31 (1) (2004).
[8] Z. Burda, A. Goerlich, A. Jarosz, J. Jurkiewicz, Signal and noise in correlation matrix, Phys. A 343 (2004) 295.
[9] L. Laloux, P. Cizeau, J.-P. Bouchaud, M. Potters, Int. J. Theor. Appl. Finance 3 (2000) 391.
[10] M. Potters, J.-P. Bouchaud, L. Laloux, Financial applications of random matrix theory: old laces and new pieces, Acta Phys. Pol. B 36 (2005) 2767.
[11] S. Pafka, I. Kondor, Noisy covariance matrices and portfolio optimization II, Phys. A 319 (2003) 487–494.
[12] S. Pafka, I. Kondor, G. Nagy, Noise sensitivity of portfolio selection under various risk measures, J. Banking Finance 31 (5) (2007) 1545–1573.
[13] R.F. Engle, Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation, Econometrica 50 (1982) 987–1008.
[14] T. Bollerslev, Generalized autoregressive conditional heteroscedasticity, J. Econometrics 31 (1986) 307–327.
[15] T. Bollerslev, R.F. Engle, J.M. Wooldridge, A capital asset pricing model with time-varying covariances, J. Political Econ. 96 (1988) 116–131.
[16] R.F. Engle, K.F. Kroner, Multivariate simultaneous generalized ARCH, Econometric Theory 11 (1) (1995) 122–150.
[17] R.F. Engle, V. Ng, M. Rothschild, Asset pricing with a factor ARCH covariance structure: empirical estimators for treasury bills, J. Econometrics 45 (1990) 213–237.
[18] T. Kariya, Mtv model and its application to the prediction of stock prices, in: T. Pullila, S. Puntanen (Eds.), Second International Tampere Conference in Statistics, University of Tampere, Finland, 1988, pp. 161–176.
[19] C. Alexander, A. Chibumba, Multivariate orthogonal factor GARCH, Technical Report, Mimeo, University of Sussex, 1997.
[20] C. Alexander, A primer on the orthogonal GARCH model, Technical Report, Mimeo, ISMA Centre, 2000.
[21] T. Bollerslev, Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model, Rev. Econ. Stat. 72 (1990) 498–505.
[22] Y.K. Tse, A.K.C. Tsui, A multivariate GARCH model with time-varying correlations, J. Bus. Econ. Stat. 20 (2002) 351–362.
[23] R.F. Engle, Dynamic conditional correlation—a simple class of multivariate GARCH models, J. Bus. Econ. Stat. 20 (2002) 339–350.
[24] R.F. Engle, K. Sheppard, Theoretical and empirical properties of dynamic conditional correlation mulrivariate GARCH, Technical Report, Mimeo, UCSD, 2001.
[25] L. Bauwens, S. Laurent, J. Rombouts, Multivariate GARCH models: a survey, CORE Discussion Paper No. 2003/31, ⟨http://ssrn.com/abstract=411062⟩, April 2003.
[26] Z. Burda, J. Jurkiewicz, M.A. Nowak, Is econophysics a solid science?, Acta Phys. Pol. B 34 (2003) 87–132.
[27] M. Mzard, A. Montanari, Constraint Satisfaction Networks in Physics and Computation, Oxford University Press, 2006, to appear. Available at ⟨http://ipnweb.in2p3.fr/~lptms/membres/mezard/main.pdf⟩.
[28] N. Gulyas, I. Kondor, to appear.
[29] V. Plerou, P. Gopikrishnan, B. Rosenow, L.A.N. Amaral, T. Guhr, H.E. Stanley, e-print cond-mat/0108023.
[30] C. Reese, B. Rosenow, Predicting multivariate volatility, cond-mat/0304082, 2003.
[31] J.D. Noh, Model for correlations in stock markets, Phys. Rev. E 61 (2000).
[32] S. Pafka, I. Kondor, Noisy covariance matrices and portfolio optimization, Eur. Phys. J. B 27 (2002) 277–280.
[33] S. Pafka, I. Kondor, Estimated correlation matrices and portfolio optimization, Phys. A 343 (2004) 623–634.